IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

VANTAGE POINT TECHNOLOGY, INC., Plaintiff, v. AMAZON.COM, INC., *et al*., Defendants.	)	Civil Action No. 2:13-cv-909-JRG **LEAD CASE** JURY TRIAL DEMANDED

DEFENDANTS' P.R. 4-5(B) RESPONSIVE CLAIM CONSTRUCTION BRIEF

TABLE OF CONTENTS

TABLE OF CONTENTS
(continued)

Page

TABLE OF AUTHORITIES

Page

CASES

TABLE OF AUTHORITIES
(continued)

Page

STATUTES

I. INTRODUCTION

One of the two patents-in-suit—U.S. Patent No. 6,374,329 (the "'329 Patent," Dkt. 192-1)—was previously litigated in this district by the original assignee of the patents-in-suit, Intergraph Corp. In the *Intergraph* litigation, Judge Davis issued a 17-page order explaining the technology of the '329 Patent and construing ten claim terms. Ex. 1. Citing no change in law or facts, Vantage Point Technology, Inc. ("VPT") asks this court to expressly contradict several of Judge Davis' constructions on issues the *Intergraph* court specifically analyzed, including Judge Davis' careful analysis of the meaning of "evicted" and his finding that the word "unknown" required no construction. Defendants propose constructions that are consistent with Judge Davis' Order, and seek only to address new issues not resolved in the prior litigation stemming from differences between the accused technology in this case and that in the *Intergraph* case.

For the second patent-in-suit—U.S. Patent No. 5,463,750 (the "'750 Patent," Dkt. 192-3)—Defendants propose constructions that are supported by the specification, the file history, and the extrinsic evidence. VPT, on the other hand, seeks broad constructions of claim terms that would far exceed the scope of the patent disclosure and contradict even the inventor's understanding of his own invention. VPT's vague constructions leave issues unresolved and will likely lead to additional disputes as the case proceeds through expert reports and trial.

II. THE '329 PATENT[1]

The '329 Patent discloses an "external-cache controller"[2] used to monitor the on-chip, internal caches of computer processors in a multi-cluster server system. *See* '329 Patent at 1:17-32. The disclosed "external-cache controller" is designed to work with commercially available,

[1] The '329 Patent is only asserted against defendants Samsung Electronics Co., Ltd., Samsung Electronics America, Inc., Samsung Telecommunications America, LLC, and MediaTek USA Inc. Thus, Section II of this brief is submitted only by the above-listed defendants.

off-the-shelf processors by tracking the cache status of each processor "from a vantage point external to the processor." *Id.* at 1:48; Ex. 2 (Intergraph Tr.) at 20:11-21:14. Claim 1 recites an "external-cache controller" that "deriv[es] the status of the private processor cache," the contents of which is "unknown externally to its associated microprocessor." '329 Patent at claim 1.

A. "private processor cache, the contents of each such private cache being unknown externally to its associated microprocessor" (claim 1)

Plaintiff's Construction	Defendants' Construction
"a cache within and dedicated to a processor whose contents cannot be read directly by any device external to that processor, including at least the external cache controller and external tag memory"	"a cache within and dedicated to a processor whose contents are unknown to any device external to that processor, including at least the external cache controller and external tag memory"

The parties dispute whether the word "unknown" should be given its plain and ordinary meaning. Defendants submit the word "unknown" requires no construction. VPT contends it should be rewritten in a way that contradicts its plain meaning and conflicts directly with the patentee's own statements. VPT's proposed construction must be rejected because it violates basic rules of claim construction and is unsupported by any evidence.

1. The Word "Unknown" Requires No Construction

"Unknown" requires no construction because—as Judge Davis correctly concluded during the *Intergraph* litigation—"its meaning is clear on its face." Ex. 1 at 10. There is no reason to deviate from the readily understood plain and ordinary meaning of "unknown" because the patentee neither "act[ed] as his own lexicographer" nor "disavow[ed] the full scope of the claim term either in the specification or during prosecution." *Hill-Rom Servs. v. Stryker Corp.*, 755 F.3d 1367, 1371-72 (Fed. Cir. 2014).

VPT relies on the lexicography exception. Dkt. 192 at 5. But that exception is

[2] The term "external cache controller" is not disputed and both parties agreed to a construction prior to VPT's opening brief. *See* Ex. 3.

inapplicable unless the patentee "clearly set[s] forth a definition" and "clearly express[es] an intent to define the term." *GE Lighting Solutions, LLC v. AgiLight, Inc.*, 750 F.3d 1304, 1309 (Fed. Cir. 2014). The intrinsic record includes neither. The '329 Patent specification never uses the word "unknown." Although used during prosecution, the patentee neither expressed an intent to define "unknown" nor equated it with "cannot be read directly." *See, e.g.*, Ex. 4 ('329 FH, 5/26/00 Resp.) at 4-6.

VPT's construction contradicts the plain meaning of "unknown" by allowing the term to encompass a design in which the external cache controller "knows … indirectly" the contents of the private cache. Dkt. 192 at 6. The meaning of "unknown" is unequivocal; it does not encompass indirect knowledge. VPT's attempt to rewrite this limitation under the guise of claim construction is impermissible. *See K-2 Corp. v. Salomon S.A.*, 191 F.3d 1356, 1364 (Fed. Cir. 1999) ("Courts do not rewrite claims; instead, we give effect to the terms chosen by the patentee."). VPT argues that its construction is supported by the "stated purpose of the invention." Dkt. 192 at 6. But the purpose of claim 1 is to use an "external-cache controller" to derive the ***status*** of private processor cache (*e.g.*, whether a cache line is in a "modified" state), the ***contents*** of which is unknown. *See* '329 Patent at claim 1. This derivation of cache status does not mean the cache content becomes known. *See* Ex. 1 (*Intergraph* Order) at 14 ("[T]he state or status of a cache line is therefore distinct from the data itself.").

VPT also argues that unless "unknown" is construed to encompass indirect knowledge, the plain meaning of the claim would exclude an embodiment in which an "optional external cache memory" stores data "inclusive" of the private cache. Dkt. 192 at 7. The scope of a claim cannot be expanded during claim construction to capture an unclaimed embodiment—especially one that is optional. In *Helmsderfer v. Bobrick Washroom Equip., Inc.*, the Federal Circuit

rejected the patent owner's proposed construction by explaining that because "[t]he patentee chooses the language and accordingly the scope of his claims," the court "cannot construe [the claims] to encompass the preferred embodiment or other illustrated embodiments" in a manner contradictory to the claims' plain meaning. 527 F.3d 1379, 1383 (Fed. Cir. 2008). Here, the '329 Patent patentee chose the word "unknown," and VPT cannot rewrite the claim during litigation to capture an optional embodiment that is not encompassed by its plain meaning.

2. VPT's Construction Contradicts the Intrinsic Evidence

As VPT admits, the patentee made an unequivocal representation to the PTO that "[t]he contents of the [private] cache ***are not known externally***, and must be tracked." Dkt. 192 at 6 (emphasis added). VPT cites no support for a construction that contradicts this clear statement made during patent prosecution. *See Typhoon Touch Techs., Inc. v. Dell, Inc.*, 659 F.3d 1376, 1381 (Fed. Cir. 2011) ("The patentee is bound by representations made and actions that were taken in order to obtain the patent.").

Moreover, VPT's construction would encompass prior art the patentee distinguished to obtain allowance. During prosecution, the patentee relied on the "unknown externally" limitation to distinguish the Fletcher prior art. *See* Ex. 4 ('329 FH, 5/26/00 Resp.) at 3-5; Ex. 5 (Fletcher). Fletcher discloses a multiprocessor system wherein each processor includes a "private CPU cache." Ex. 5 at 2:39-42. The contents of Fletcher's private cache are not directly accessible, but the private cache directory is replicated in copy directories ("CDs") in an external Storage Controller ("SC"). *Id.* at 5:65-66, 6:17-23. When another processor makes a cache request, the external copy directory is interrogated rather than the private processor cache. *Id.* at 9:9-24. Thus, Fletcher's private processor cache stores data not directly accessible by external devices, but that are externally known through copy directories in the external storage controller.

To gain allowance over Fletcher, the patentee argued that the prior art failed to disclose

the "unknown externally" limitation. *See* Ex. 4 ('329 FH, 5/26/00 Resp.) at 4-5. According to the patentee, the contents of Fletcher's private cache "are not hidden" because "[a] copy of each processor's cache directory is provided to an external storage controller." *Id.* at 4. In other words, the patentee argued that the contents of Fletcher's private cache are known to the external storage controller and are indirectly accessible to other devices through the copy directories in the external storage controller. *See id.* at 4-5. By distinguishing Fletcher, the patentee confirmed that the plain meaning of "unknown" excludes the "inclusive" cache embodiment cited in VPT's brief (Dkt. 192 at 7).[3] Thus, VPT's proposed construction of "cannot be read directly"—which permits indirect access to a replicate copy of cache contents in an external cache memory—would improperly broaden the term to encompass the distinguished Fletcher design.[4] VPT's construction should be rejected because it violates the principle that "[p]atent claims may not be construed one way in order to obtain their allowance and in a different way against accused infringers." *Southwall Tech., Inc. v. Cardinal IG Co.*, 54 F.3d 1570, 1576 (Fed. Cir. 1995).

B. "external tag memory" (claim 1)

Plaintiff's Construction	Defendants' Construction
"Storage that is situated apart from the main memory that includes information about the status of the private processor cache."	"storage that is situated apart from the processor and main memory that includes information about the status of the private processor cache and which has a larger number of entries than the number of cache lines in the

[3] Although the patentee also distinguished Fletcher on other grounds, its reliance on the "unknown" limitation is independently binding. *See Andersen Corp. v. Fiber Composites, LLC*, 474 F.3d 1361, 1374 (Fed. Cir. 2007) ("An applicant's invocation of multiple grounds for distinguishing a prior art reference does not immunize each of them from being used to construe the claim language. Rather, as we have made clear, an applicant's argument that a prior art reference is distinguishable on a particular ground can serve as a disclaimer of claim scope even if the applicant distinguishes the reference on other grounds as well.").

[4] Indeed, the patentee originally drafted a claim that captures the idea of a private cache that cannot be read directly, but abandoned it during prosecution. *See* Ex. 6 ('329 FH, 2/4/99 Prel. Am.) at 4 (claim 32 reciting "wherein each processor has an internal cache which is externally unreadable outside the processor"); Ex. 7 ('329 FH, 12/22/99 Resp.) at 1 (abandoning claim 32).

	private processor caches associated with the first processor cluster"

The dispute for this term involves an issue not raised during the *Intergraph* litigation. In that prior case, the parties disagreed over whether the word "external" requires the tag memory to be external to the processor. *See* Ex. 1 at 10-11. The parties in this case adopt Judge Davis' construction of "external." Defendants' proposed construction further clarifies the scope of the limitation by describing the ***only*** manner in which the claimed invention can possibly operate.

As VPT admits, "the invention [of the '329 Patent] ***needs to forcibly track*** the L2 cache's contents. Towards this end, the XAP 204 tracks evicted P6 cache lines in producing a correct reflection of the P6 internal caches." Dkt. 192 at 11 (citing '329 Patent at 7:31-34) (emphasis added). The claimed invention "needs" to perform the disclosed tracking because its goal is to work with a processor design that does not "disclose[] how to make public the contents of the level 2 (L2) cache serving a processor cluster." '329 Patent at 7:19-21.

The ***only*** way for the disclosed invention to "forcibly track" the cache contents is to have an external tag memory that has a larger number of entries than the private processor cache. *See id.* at 15:27-30 (disclosing that external tag memory is "able to accommodate up to 16 to 32 the number of cache lines that may be stored by a single quad-CPU processor segment"). Because the claimed invention works with processors that do not disclose the contents of their private caches, the external cache controller must "derive" the status of those private caches. *Id.* at 7:19-21; claim 1. It does so by keeping track of what ***might*** be in the private cache via tracking data entering and exiting the processor and, then, by "snooping" to eliminate those items that are no longer in the cache. *Id.* at 15:42-44, 15:51-60, claim 1. Because snooping cannot occur after every moment that the cache contents may change, the external tag memory must have capacity to store statuses of items that ***might*** be in each cache line. *Id.* at 7:14-18, 15:45-60. Thus, to

perform the "forcibly track" function, the external tag memory must have more entries than the private processor cache. The disclosed invention cannot function otherwise.

C. "external tag memory non-hardwired to the private processor cache" (claim 1)

Plaintiff's Construction	Defendants' Construction
"The private processor cache does not duplicate its status in the external tag memory or the private processor cache does not use the same logical address as the external tag memory."	"external tag memory not permanently connected to the private processor cache such that, for example, the private processor cache does not duplicate its status in the external tag memory or the private processor cache does not use the same address as the external memory"

The parties disagree whether certain statements made by the patentee during prosecution constitute a definition (VPT's position) or a disclaimer (Defendants' position). If there was a disclaimer, the parties also disagree over the scope of that disclaimer.

VPT urges the Court to adopt the construction from the *Intergraph* case. But that prior construction was undisputed and Judge Davis did not address the issues presented by the parties here. *See* Ex. 1 (*Intergraph* Order) at 11 ("The parties also seem to agree on the substantive meaning of 'non-hardwired' …."); Ex. 2 (*Intergraph* Tr.) at 121:17-124:22. The Court should "place[] little weight on that earlier agreement." *Civix-DDI, LLC v. Hotels.com, L.P.*, 2010 WL 4386475, at *3 (N.D. Ill. Oct. 24, 2010); *see also, Fuji Photo Film Co. v. ITC*, 386 F.3d 1095, 1101 (Fed. Cir. 2004) (ALJ's acceptance of stipulated construction does not "constitute a formal claim construction"). As explained below, the patentee's prosecution history statements do not constitute a definition of the disputed term, but instead serve as a disclaimer of claim scope.

1. The Patentee Did Not Act As Its Own Lexicographer

The term "non-hardwired" does not appear in the specification, and the patentee never defined this term in the intrinsic record. VPT relies solely on the following passage from the prosecution history to support its position that the patentee defined "non-hardwired":

> In contrast, Fletcher discloses a storage controller that includes copy directories at col. 6, lines 19 and 20, referencing Flusche et al. The copy directories duplicate and have the same logical address as each processor's cache directory, as described by Flusche at col. 5, lines 39-44. Hence, in Fletcher, the copy directories are ***hardwired*** to the processor cache, unlike claim 26, which requires that the external tag memory be ***non-hardwired*** to the private processor cache.

Dkt. 192 at 9 (citing Ex. 8 at 7) (emphasis added).

The cited passage neither "clearly set[s] forth a definition" for "non-hardwired" nor "clearly express[es] an intent to define the term." *GE Lighting Solutions*, 750 F.3d at 1309. Rather, it assumes that the Examiner understood the common meaning of "hardwired" and its characteristics. The Patentee argued that the description of the Fletcher/Flusche design—its external and internal memories store the same data and use the same address—reflects the characteristics of a permanently connected or "hardwired" design. Nothing in the cited passage suggests that "hardwired" is being used in a special way that contradicts its plain meaning. "[T]o deviate from the plain and ordinary meaning of a claim term to one of skill in the art, the patentee must, with some language, indicate a clear intent to do so in the patent." *Hill-Rom Servs.*, 755 F.3d at 1373. VPT fails to demonstrate any evidence of such intent.

Because the lexicography exception is inapplicable, the claim construction analysis must begin with the plain and ordinary meaning of the term. "Hardwired" is a common term used to describe permanent electrical connections, as opposed to temporary ones.[5] A tag memory non-hardwired to the private processor cache thus cannot be permanently connected to that cache.

[5] *See e.g.*, Ex. 9 WEBSTER'S II NEW RIVERSIDE UNIVERSITY DICTIONARY at 505 (1995) ("Of, pertaining to, or effected by means of logic circuitry that is permanently connected within a computer or calculator."); Ex. 10 OXFORD ONLINE DICTIONARY (AMERICAN ENGLISH) ("Involving or achieved by permanently connected circuits."); Ex. 11 MERRIAM-WEBSTER ONLINE DICTIONARY ("implemented in the form of permanent electronic circuits"); Ex. 12 Arora, FOUNDATIONS OF COMPUTER SCIENCE at 42 (2006) ("Some personal computers have cache memory chips hardwired onto the motherboard."); Ex. 13 Lombardi, COMPUTER LITERACY 100 (1983) ("Hard wired also refers to changes and modifications that are made to computer

VPT does not appear to dispute that Defendants' construction accurately reflects the plain meaning of "non-hardwired," but argues that this meaning is inconsistent with the specification. *See* Dkt. 192 at 9-10. VPT's argument is baseless. The '329 Patent is directed to providing an external cache controller that would allow any commercially-available, off-the-shelf processor to be used in multi-cluster systems. '329 Patent 1:27-32; Ex. 2 (*Intergraph* Tr.) at 20:11-21:14. All processors of the preferred embodiments are detachably connected to a motherboard and other processors through sockets that allow users to easily change processors.[6] For example, the Intel Pentium Pro Developer's Guide—incorporated by reference and thus part of the specification—discloses that the P6 processors connect to other components through removable "Zero Insertion Force" sockets to allow for a flexible number of processors. *See* '329 Patent at 3:26-35; Ex. 14 at 11-8 to 11-10, 17-1 to 17-6. VPT fails to cite any evidence showing that any disclosed processor is "soldered" to external tag memory. *See* Dkt. 192 at 9-10.

2. The Patentee Disclaimed Coverage of Any System Where the External Tag Memory Duplicates the Status and Uses the Same Addresses

Rather than a definition, the patentee's arguments distinguishing the Fletcher/Flusche design serve as a disclaimer of claim scope. *See Andersen Corp.*, 474 F.3d at 1374 ("an applicant's argument that a prior art reference is distinguishable on a particular ground can serve as a disclaimer of claim scope"); *Omega Eng'g, Inc. v. Raytek Corp.*, 334 F.3d 1314, 1324 (Fed. Cir. 2003) (prosecution disclaimer "narrows the ordinary meaning of the claim congruent with the scope of the surrender"). The patentee admitted the prior art "discloses a storage controller that includes copy directories" that "duplicate and have the same logical address as each

---

equipment that actually permanently change the machine as opposed to plug in changes with accessory boards.")

[6] *See* '329 Patent at 1:27-32; Ex. 15 (Pentium Pro DataSheet) at 1; Ex. 16 (Cyrix 6x86 Datasheet) at 2; Ex. 17 (Cyrix MII Databook) at 3; Ex. 18 (AMD K6 Thermal Note) at 4 (Fig. 3); Ex. 19 (AMD K5 Data sheet) at 1.

processor's cache directory." Ex. 8 ('329 FH, 9/7/01 Resp.) at 7; Ex. 20 ('329 FH, 11/15/00 Resp.) at 8. The parties disagree whether this disclaimer applies only to a logically-addressed design, or is agnostic to the address scheme. Flusche notes that the type of address used by its cache directory can be virtual or real. *See* Ex. 21 (Flusche) at 11:22-24 ("An IE command to the BCE (Fig. 3A) may contain the logical address (i.e. ***virtual or real address***) for a DW instruction or operand of data.") (emphasis added). Thus, the scope of "non-hardwired" cannot include any design where the external tag memory duplicates the cache status and uses the same address, regardless of the type of address scheme used.

D. "deriving the status of the private processor cache associated with the first processor cluster by tracking evicted cache lines and data entering and exiting the first processor cluster" (claim 1)

Plaintiff's Construction	Defendants' Construction
"to get or obtain the status of the private processor cache associated with the first processor cluster by tracking: (1) all modified lines within the private processor cache of that same cluster, (2) all data that enters that same cluster, and (3) all data that exits that same cluster, and using at least that information to determine the status of each private processor cache in only that same cluster."	"to get or obtain the status of the private processor cache associated with the first processor cluster by tracking: (1) all modified lines within the private processor cache of that same cluster, (2) all data that enters that same cluster, and (3) all data that exits that same cluster, and using at least that information to determine the status of each private processor cache in only that same cluster. 'Evicted' does not mean what was in the cache but is no longer there (i.e., a line that has been ejected), but rather something that is now in the cache that is different from what was previously there (i.e., a modified line)"

VPT concedes that it must convince the Court to reject and overturn Judge Davis' extensive and detailed analysis of this term from the *Intergraph* case, but VPT does not address the substance of that analysis. *See* Dkt. 192 at 11. VPT cites to the exact same passages that Judge Davis quoted and discussed. *Compare id. with* Ex. 1 (*Intergraph* Order) at 13. VPT advocates for an opposite conclusion, but does not explain (a) how the same passages analyzed by Judge Davis support VPT's position or (b) why Judge Davis was wrong. Defendants do not needlessly repeat Judge Davis' analysis herein, but submit that Judge Davis was correct for the

reasons explained in that decision. *See* Ex. 1 at 12-15; Ex. 2 at 99-104.

VPT repeats the ***exact*** same argument that Judge Davis considered and rejected. In the prior case, "Intergraph argue[d] that [evicted] means 'tracking the replacement of unmodified cache lines.'" Ex. 1 at 13. In this case, VPT argues that "evicted" include "tracking of unmodified cache lines." Dkt. 192 at 11. Indeed, VPT even relies on the same passage cited by Intergraph, and its brief reprises the exact arguments that Intergraph presented to Judge Davis during oral arguments. *See id.* at 11; Ex. 1 at 12-13; Ex. 2 at 94-106.

Judge Davis conducted a careful analysis of the prosecution history and specification before concluding that "[r]ather than referring to what was in the cache but is no longer there … ["evicted"] refers to something that is now in the cache that is different from what was previously there (*i.e.*, a modified line)." Ex. 1 at 13. Judge Davis' analysis confirms that VPT's position is incorrect and that "'evicted cache lines' cannot refer to unmodified data exiting the processor cluster." *Id*. at 14. While VPT quotes a specification passage cited by Judge Davis, VPT fails to explain why Judge Davis' conclusion was wrong. VPT also quotes a part of the specification not cited by Judge Davis, but does not explain its relevance. The fact that a "Bus Read and Invalidate (BRIL) and a Bus Invalidate Line (BIL) function may or may not modify the requested line" (Dkt. 192 at 11) does not convey any information about whether tracking "evicted" cache lines is properly construed as tracking modified lines. In sum, VPT fails to explain how any evidence supports its position.

Because VPT disputes the very basis of the prior construction, Defendants' proposed construction incorporates Judge Davis' explicit definition of "evicted" to clarify the scope of the claim limitation at issue. *See* Ex. 1 at 13. This definition of "evicted" was the heart of Judge Davis' construction for the disputed claim term. Its inclusion would avoid jury confusion by

expressly defining an important word within the disputed claim term.

VPT's insistence on leaving ambiguity in the "deriving …" limitation by not defining "evicted" is unsurprising given the enormous and dispositive difference between the accused products and the '329 Patent. The passage VPT cites highlights this fact:

> A problem that a preferred embodiment addresses is that no Intel specification for the P6 processor discloses how to make public the contents of the level 2 (L2) cache serving a processor cluster. No method is disclosed for determining 30 which internal cache line is being replaced when a new one is fetched. … Consequently, the invention needs to forcibly track the L2 cache's contents. Towards this end, the XAP 204 tracks evicted P6 cache lines in producing a correct reflection of the P6 internal caches.

'329 Patent at 7:18-34. The accused processors in this case operate in exactly the opposite manner: the accused processors ***do*** disclose "which internal cache line is being replaced when a new one is fetched." There is a separate signal which explicitly tells the cache controller exactly which cache line is being replaced. Thus, the accused products have no need to derive the status of the processor cache status by "forcibly track[ing] the L2 cache's contents." This derivation is unnecessary because the accused system knows exactly what occurs in the processor through notifications from the processor cache. This reflects the difference between a tightly integrated system in which both the processor and the cache controller are supplied by a single entity (ARM) and the patented system in which Intel supplied a processor and Intergraph developed an external controller to track and infer the cache status within the Intel processor.

E. "tag controller" (claim 1)

Plaintiff's Construction	Defendants' Construction
"a controller that manages the external tag memory"	"a controller that manages the external tag memory for processors that do not disclose which internal cache line is being replaced when a new one is fetched"

VPT admits that the claimed invention works on a system in which "no Intel specification for the P6 processor discloses how to make public the contents of the level 2 (L2) cache serving

a processor cluster. No method is disclosed for determining 30 which internal cache line is being replaced when a new one is fetched." Dkt. 192 at 11 (citing '329 Patent at 7:18-34). For this reason, the "***invention*** needs to forcibly track the L2 cache's contents." *Id*. (emphasis added). Defendants' proposed construction for "tag controller" describes this fundamental "need" of the invention. Indeed, VPT concedes that the invention is applicable to systems in which the "processor hides such [internal cache line] information." *Id.* at 13.

Defendants' construction of "tag controller" accurately reflects the scope of the invention as described by the specification. *See Regents of the Univ. of Minn. v. AGA Med. Corp*., 717 F.3d 929, 936 (Fed. Cir. 2013) ("When a patent … describes the features of the 'present invention' as a whole, this description limits the scope of the invention."). Defendants' construction does not conflict with Judge Davis's construction of "tag controller associated with the first processor cluster" because the dispute in the *Intergraph* case centered on the meaning of the word "associated" and not on the scope of "tag controller." *See* Ex. 1 at 12.

F. <u>"when the processor bus is idling" (claim 2)</u>

Plaintiff's Construction	Defendants' Construction
"when the processor bus is not being used by the processor"	"when the processor bus is not being used by the processor or other connected device"

Defendant's construction reflects the plain and ordinary meaning of this term and is supported by the intrinsic evidence. VPT did not address this term in its brief, and its construction contradicts the intrinsic record.

A processor bus can connect several devices in a computer system such as the processor, memory, and input/output controllers. *See e.g.*, Ex. 14 at 1-3; '329 Patent at 3:26-35 (incorporating Ex. 14 by reference). To be idle, the bus cannot be used by any connected device. *See* '329 Patent 7:14-18 (noting there must be "available bus cycles" for bus to be idle). VPT's

proposal would allow a processor bus actively being used by non-processor components to be considered "idle." VPT's construction not only contradicts the plain meaning of the term, but would rewrite the limitation from the bus being idle to the processor being idle. This change would also allow VPT to recapture claim scope distinguished during prosecution. *See* Ex. 20 ('329 FH, 11/15/00 Resp.) at 4, 6 (patentee amending claim from the processor being idle to the bus being idle); Ex. 22 ('329 FH, 8/16/00 OA) at 4, 7-8 (Examiner finding obvious snooping while the processor is idle and that application lacked support for the same). Thus, the processor bus is "idling" when it is not used by the processor and any other connected devices.

G. "to identify if the data has been modified if the tag controller indicates that data is held within the private processor cache associated with the first processor cluster" (claim 2)

Plaintiff's Construction	Defendants' Construction
Plain and ordinary meaning	"asking the private processor cache if particular data is present to determine whether a cache line has been modified if the tag controller indicates a copy is loaded in one of the first processor cluster's private processor caches"

Defendants' construction reflects the plain and ordinary meaning of this term—as VPT acknowledges in its briefing. *See* Dkt. 192 at 13. Claim 2 recites "a snooper for snooping the private processor cache of the at least one microprocessor associated with the first processor cluster" before the "to identify" clause. *See* '329 Patent at claim 2. In the context of the claim language, it is clear that the "snooping" function requires making an inquiry into the contents of the private processor cache. Defendants' construction is also fully supported by the requirement in claim 2 that "snooping" occurs "when the processor bus is idling." *See id.* Defendants' construction clarifies claim scope and puts the technical language of the claim limitation in a more jury-friendly form. Because VPT does not disagree with the substance of Defendants' clarification of the claim language, Defendants' construction should be adopted.

H. "a snooper for snooping the private processor cache of the at least one microprocessor associated with the first processor cluster to identify if the data has been modified if the tag controller indicates that data is held within the private processor cache associated with the first processor cluster" (claim 2)

Plaintiff's Construction	Defendants' Construction
Plain and ordinary meaning (not subject to 112 ¶ 6)	Governed by 35 U.S.C. § 112 ¶6. **Function**: snooping the private processor cache of the at least one microprocessor associated with the first processor cluster to identify if the data has been modified if the tag controller indicates that data is held within the private processor cache associated with the first processor cluster **Structure**: There is insufficient disclosure of structure to perform this function. Alternatively, the closest corresponding structure is XAP 204 in Figure 9 as referenced at 4:51-65.

The words of the ***claim*** must convey sufficient structure to avoid a construction under 35 U.S.C. § 112(6). *See Robert Bosch, LLC v. Snap-On Inc.*, 769 F.3d 1094, 1099 (Fed. Cir. 2014) ("we ask if the claim language, read in light of the specification, recites sufficiently definite structure to avoid § 112, ¶ 6"); *Welker Bearing Co., v. PHD, Inc*., 550 F.3d 1090, 1095 (Fed. Cir. 2008) ("purely functional limitations that do not provide the structure that performs the recited function" should be construed under § 112(6)). VPT does not identify a single term in the "snooper" limitation that is structural. The phrase "snooper for snooping" is written as a purely functional term. Indeed, a "snooper" describes only the function of this limitation (*i.e.*, for snooping) and is much less structural than the "program recognition device" and "program loading device" held to be governed by § 112(6) in *Bosch*. *Robert Bosch*, 769 F.3d at 1099. Thus, "snooper for snooping" should be governed by § 112(6).

A term governed by § 112(6) is invalid if the patent specification does not disclose sufficient structure corresponding to the recited function. *See Robert Bosch*, 769 F.3d at 1102. VPT points to passages in the specification that describes the snooping function, but fails to identify any structure corresponding to that function. *See* Dkt. 192 at 14. Figure 9 of the '329

Patent illustrates "snooper" as simply an oval within the XAP block. *See* '329 Patent at 4:51-65, Fig. 9. A "black box" illustration cannot provide sufficient structural disclosure for a term governed by § 112(6). *See ePlus, Inc. v. Lawson Software, Inc*., 700 F.3d 509, 518 (Fed. Cir. 2012). It is also not sufficient that one skilled in the art might be able to design a "snooper"—rather, the specification must clearly disclose this structure, and clearly link such disclosed structure to the recited function. *See id.* at 519 (rejecting argument that the structure of a § 112(6) term was known in the art). Because the "snooper" limitation is purely functional and the specification fails to disclose any structure linked to its recited snooping function, claim 2 is indefinite and therefore invalid. *See Robert Bosch*, 769 F.3d at 1102.

I. "processor cluster" (claim 1)

Plaintiff's Construction	Defendants' Construction
"cluster having at least one of each of a microprocessor, a processor bus, and a private processor cache"	"a set of one or more processors in a multiprocessor system that share a common CPU bus"

VPT's construction—based entirely on importing other claim limitations into the construction of this term—eviscerates "processor cluster" as a limitation and should be rejected. *See Innova/Pure Water, Inc. v. Safari Water Filtration Sys*., 381 F.3d 1111, 1119 (Fed. Cir. 2004) ("all claim terms are presumed to have meaning in a claim"); *Bicon, Inc. v. Straumann Co.*, 441 F.3d 945, 950 (Fed. Cir. 2006) ("[C]laims are interpreted with an eye toward giving effect to all terms in the claim."). The claim already includes nearly the exact language with which VPT is attempting to define "processor cluster." *See* '329 Patent at claim 1 ("each cluster having at least one microprocessor having a processor bus and a private processor cache"). Thus, VPT's construction renders "processor cluster" entirely superfluous.

Defendants' construction, on the other hand, gives meaning to "processor cluster" and is supported by the patent specification. The specification defines the "present invention" as a

multiprocessor system. *See* '329 Patent at 1:17-19, 1:33-37; *Regents of the Univ. of Minn.*, 717 F.3d at 936 ("When a patent … describes the features of the 'present invention' as a whole, this description limits the scope of the invention."). The specification further defines a "processor cluster" by the common bus shared by each group of processors. *See* '329 Patent at 1:33-40. As the specification notes, the P6 processors of the preferred embodiment can only be grouped in clusters of up to four processors because the CPU bus allows "a maximum of 16 outstanding transactions," and each processor can "post … four transactions to the P6 bus." *See id.* at 1:57-59, 3:18-20, and 6:9-15. Thus, the maximum number of processors allowed in each cluster is constrained by the capacity of their shared bus. VPT's construction ignores this important consideration and would allow any arbitrary group of processors to be considered a "processor cluster" even when they do not share a common bus.[7]

III. THE '750 PATENT

The '750 Patent relates to the use of translation lookaside buffers ("TLBs") to improve the performance of virtual-to-physical memory address translation. *See* '750 Patent at 1:9-12. As the patentee admitted during prosecution, TLBs and direct address translators were known in the art. Ex. 23 ('750 FH, 7/28/94 OA) at 2-3; Ex. 24 ('750 FH, 9/29/94 Resp.) at 1-2. In fact, the patentee relied on an incorporated prior art patent to supply the required technical disclosure for the claimed invention. *Id.*; '750 Patent at 3:16-40 (incorporating U.S. Patent No. 4,933,835 by reference). What the '750 Patent discloses is a narrow improvement on a multi-pipeline computer system that adds a separate TLB to each instruction pipeline. '750 Patent at 4:11-15.[8]

[7] VPT's argument against the use of "CPU bus" in Defendants' construction is baseless because the specification uses the terms "CPU" and "processor" interchangeably. *See, e.g.*, '329 Patent at 1:33-37, 3:15-17.

[8] *See also* Ex. 25 ('750 FH, 2/28/95 Resp.) at 8 ("The claimed invention is directed to a method and apparatus for translating virtual address[es] in a computing system having multiple pipelines wherein a separate translation lookaside buffer (TLB) is provided for each pipeline.").

As defined by the specification, this computer system of the "Present Invention" includes multiple instruction pipelines that share system resources, including a common "instruction issuing unit 14" and a common "dynamic translation unit" ("DTU 162"). *See id.* at 1:9-34, Figs. 1 and 5. The '750 Patent does not cover systems having multiple, independently operating processors—*i.e.*, instruction pipelines that do not share resources—because these systems were distinguished during prosecution. *See, e.g.*, Ex. 24 ('750 FH, 9/29/94 Resp.) at 3 (distinguishing "a system with multiple cache memories for multiple processors").

A. "[first/second] instruction pipeline" (claim 8)

Plaintiff's Construction	Defendants' Construction
"[first/second] of at least two structures, each consisting of a sequence of stages"	"[first/second] of at least two structures, each consisting of a sequence of stages that execute instructions received from a single instruction issuing unit"

The parties have two disputes regarding the "[first/second] instruction pipeline" term: (1) whether an "instruction pipeline" must execute instructions, and (2) whether the two instruction pipelines must receive instructions from the same instruction issuing unit. Defendants' construction answers both questions in the affirmative and is supported by the intrinsic and extrinsic evidence. VPT's construction to the contrary is incomplete, overbroad, and fails to give meaning to all terms.

1. An "Instruction Pipeline" Must Execute Instructions

Logically, an ***instruction*** pipeline must execute instructions. VPT's construction fails to address this express requirement of the claim term. Many structures "consist[] of a sequence of stages" (*e.g.*, the Saturn V rocket has three sequential stages), but not every multi-stage structure is an "instruction pipeline." By excluding the requirement of executing instructions, VPT's construction would allow any "structure," including a single instruction pipeline, to be arbitrarily divided into multiple segments to meet the "first" and "second" pipeline limitation, even though

the divided segments cannot each execute instructions. VPT's attempt to read "instruction" out of this limitation fails to "giv[e] effect to all terms in the claim" and should be rejected. *Bicon, Inc.*, 441 F.3d at 950.

The intrinsic evidence confirms Defendants' construction. In a proceeding before the PTO, VPT admitted that instruction pipelines must execute instructions: "The CPUs employ pipelining, in which instructions are partitioned into various execution phases to allow simultaneous operations on different phases of different instructions by different logical units." Ex. 26 (VPT's IPR Resp.) at 5.

Defendants' construction is also supported by the extrinsic evidence. Hennessy—the leading textbook on processor design—explains an instruction pipeline as follows:

> A pipeline is like an assembly line: ***Each step in the pipeline completes a part of the instruction.*** As in a car assembly line, the work to be done in an instruction is broken into smaller pieces, each of which takes a fraction of the time needed to complete the entire instruction. Each of these steps is called a pipe stage or a pipe segment. The stages are connected one to the next to form a pipe--***instructions enter at one end, are processed through the stages, and exist at the other end***.

Ex. 27 at 251 (emphasis added). Indeed, VPT's technology tutorial acknowledges that this is a correct description of an instruction pipeline. Ex. 28 (VPT's Tech. Tutorial) at 14. Thus, an instruction pipeline must include a sequence of stages that execute instructions; a multi-stage pipeline segment incapable of executing instructions is not an "instruction pipeline."

Without requiring each pipeline to execute instructions, VPT's construction would broaden the claim to cover two multi-stage units—an instruction unit and a data unit, for example—that are segments within the same instruction pipeline, so long as each unit has its own TLB. Such a design was expressly distinguished during prosecution. The inventor's own prior art patent includes separate cache memory management units (CAMMUs) for instructions and data, with each CAMMU having its own TLB. *See* Ex. 29 ('835 Patent) at Figs. 1, 2, and 8.

Because the instruction unit and the data unit are part of the same instruction pipeline, the Examiner found that the prior art does not disclose "that multiple TLBs are used for multiple pipelines." *See id.* at Fig. 2; Ex. 30 ('750 FH, 11/3/94 OA) at 3. VPT's construction would allow the instruction and data units to be considered separate "instruction pipelines," which would broaden the claim scope to cover the distinguished prior art design.

2. The First and Second Instruction Pipelines Must Receive Instructions from the Same Instruction Issuing Unit

The '750 Patent discloses a system that "includes an instruction issuing unit 14 which communicates instructions to a plurality of (e.g., eight) instruction pipelines 18A-H over a communication path 22." *Id.* at 1:14-17. The only disclosed embodiment includes multiple instruction pipelines receiving instructions from the same instruction issuing unit. *See id.* at Fig. 1; 4:59-63. The '750 Patent does not disclose, and does not claim, two unrelated and uncoordinated instruction pipelines that operate independently. Accordingly, Defendants' proposed construction includes the phrase "received from a single instruction issuing unit."

VPT fails to cite any support for its position that the two claimed instruction pipelines can receive instructions from different issuing units because there is no support for such a design. The only evidence cited by VPT, that the three pipelines illustrated in Figure 5 "may be three of the pipelines 18A-H shown in FIG. 1," actually supports Defendants' position. Dkt. 192 at 16 (citing '750 Patent at 4:62-63). All eight pipelines in Figure 1 receive instructions from the same instruction issuing unit 14, as urged by Defendants. *See* '750 Patent at 4:56-63; Figs. 1 and 5. The '750 Patent does not disclose a computer having multiple instruction issuing units. Indeed, removing the requirement of a common instruction issuing unit would allow this limitation to cover a system having two processors that operate independently without sharing a common instruction source—a design that the patentee distinguished during prosecution. *See* Ex. 24

('750 FH, 9/29/94 Resp.) at 3.

The testimony of the inventor, Mr. Howard Sachs, confirms Defendants' construction. Mr. Sachs explained that the claimed invention "assumes that the pipelines … all are running on the same thread" of computer instructions. Ex. 31 (Sachs Decl.) ¶ 7. In other words, the claimed invention requires multiple instruction pipelines to execute instructions from a common instruction stream, issued by a common instruction issuing unit as illustrated in Figure 1. Mr. Sachs further explained that "[i]f each pipeline is executing a different thread then this type of system is not comprehended by U.S. Patent No. 5,463,750." *Id.*

The proper construction of a term "can only be determined and confirmed with a full understanding of what the inventors actually invented and intended to envelop with the claim." *Medrad, Inc. v. MRI Devices Corp.*, 401 F.3d 1313, 1319 (Fed. Cir. 2005); *see also Phillips v. AWH Corp.*, 415 F.3d 1303, 1317 (Fed. Cir. 2005) (identifying "inventor testimony" as relevant extrinsic evidence); *Voice Tech. Grp., Inc. v. VMC Sys., Inc.*, 164 F.3d 605, 615 (Fed. Cir. 1999) ("An inventor is a competent witness to explain the invention and what was intended to be conveyed by the specification and covered by the claims."). Here, the inventor's testimony sheds important light on the scope of the invention, and confirms that the "[first/second] instruction pipeline" must receive instructions from a common instruction issuing unit.

B. "master translation memory" (claim 8)

Plaintiff's Construction	Defendants' Construction
Plain and ordinary meaning	"Page tables in main memory"

VPT's conclusory argument that "master translation memory" should have plain and ordinary meaning is without merit. "Master translation memory" is a coined term that does not appear in the specification, and it has a specific technical meaning in the context of claim 8. VPT's construction does not help the jury understand the scope of this claim term, and would

require the jury to address the legal question of claim interpretation. Defendants' construction should be adopted because, in the context of the invention, it is clear that the "***master*** translation memory" is the "page tables in main memory." *See* '750 Patent at 3:40-41.

Defendants' construction is supported by the claim language and the specification. Claim 8 imposes at least two requirements: (1) the existing translation data stored in the first and second translation buffers is a "subset" of the translation data in the "master translation memory," and (2) when a desired translation is not present in the translation buffers (*i.e.*, when there is a TLB miss), the new translation data is stored from the "master translation memory" into the translation buffers. *See* '750 Patent at claim 8. Claim 8 expressly calls for "storing the translation data for the [first/second] virtual address from the master translation memory into the [first/second] translation buffer." *Id.* Thus, claim 8 requires the "translation data" to be sourced from the "master translation memory" when that "translation data" is not present in the "translation buffer." *See id.* at 3:33-47. That source—the "master translation memory"—must be the master source of all "translation data" for the "translation buffer."

The master source of the translation data is the page tables in main memory. The specification explains that there are two repositories of translation memory: a TLB and page tables in main memory. When a particular translation data is not present in the TLB (*i.e.*, the "translation buffer"), then the system uses a "dynamic translation unit (DTU) 162 for accessing page tables in main memory" and storing the updated translation data from those page tables into the TLB. *Id.* at 2:67-3:1; *see also id.* at 3:40-47 ("DTU 162 accesses the page tables in main memory 34 … TLB 158 is updated through a communication path 196, and instruction issuing resumes."); 5:42-44 ("[T]he address translation tables in main memory are accessed to obtain address translation information for [storing in] any other TLB."). Thus, the TLB stores only a

subset of the translation data, while the master set of the translation data is stored in the page tables of main memory.

The figures of the '750 Patent, similarly, confirm that the TLBs are updated with translation data stored in page tables in main memory. Figure 4 shows a TLB 158 (translation buffer) that is connected to a DTU shown as transferring data to and from "main memory." Figure 5 shows three TLB's (222A-C) connected to a DTU 162 that similarly accesses "main memory" to get updated data. *See* '750 Patent at Figs. 4-5.

Thus, translation buffers (TLBs) cannot be the "***master*** translation memory" because they store only a subset of the translation data. *See id*. at 5:27-30, 5:52-55, 2:59-3:15. Instead, the TLBs are updated from the "page tables in main memory." *See id.* To give meaning to the word "master," the "master translation memory" must mean those page tables stored in main memory.

C. "direct address translation unit" (claim 8)

Plaintiff's Construction	Defendants' Construction
"a common unit shared by all instruction execution pipelines that translates a virtual memory address and also comprising the master translation memory"	"a common unit shared by all instruction execution pipelines that translates a virtual memory address using data in the page table" and also comprising the [master translation memory]"

The parties agree that the "direct address translation unit" ("DATU") must be "common" and "shared" by all instruction execution pipelines. This construction is necessary because, in a proceeding before the PTO, VPT disclaimed designs having multiple pipelines that "do not share any 'direct address translation unit.'" *See* Ex. 26 (VPT's IPR Resp.) at 6-7.

The difference between the parties' constructions has a significant impact on this case. The disputed issue is ***what*** must be common and shared between the pipelines. VPT's brief does not identify exactly what VPT believes is "common" and "shared" by the pipelines. The "unit" that is common and shared cannot be ***any*** processing unit. Rather, it must be the unit that

performs "direct address translation." The only unit that performs such "direct address translation" is the circuitry (or software) that uses data in the page tables to perform address translations—converting virtual addresses into physical addresses. Defendants' construction clarifies this issue and identifies an operable structure, while VPT's construction seeks to avoid and defer the issue.

VPT should not be permitted to avoid addressing this issue because, if left unresolved at the present stage, the question of what is common and shared will undoubtedly arise during expert reports, during pre-trial proceedings, and potentially at trial. Defendants respectfully submit that the Court should adopt Defendants' construction because it defines the meaning of "direct address translation" and provides the needed resolution to a disputed term.

1. A Direct Address Translation Must Use Data from Page Table

Defendants' construction correctly excludes indirect translations from the scope of "direct address translation." The '750 Patent discloses two ways to obtain a physical address from a virtual address: (1) retrieve pre-translated address data buffered in memory (*i.e.*, from a TLB); and (2) when there is no pre-translated address (*i.e.*, a TLB miss), use a translation unit to "access[] page tables in main memory" to perform the translation. '750 Patent at 2:58-3:6; 5:2-25. Only the second method is a "***direct*** address translation." *See id.* at Abstract ("activating the direct address translation circuit when the translation data for the virtual address is not stored in the TLB").

VPT does not, and cannot, identify any support for its position that a direct address translation can be performed without using a page table. In fact, the PTO found that the '750 Patent specification fails to sufficiently explain any direct address translation. *See* Ex. 23 ('750 FH, 7/28/94 OA) at 2-3 (rejecting all claims because the specification lacks sufficient disclosure of direct address translation). During prosecution, the patentee relied on the incorporated '835

Patent to support its claim of "address translators." Ex. 24 ('750 FH, 9/29/94 Resp.) at 1-2. The '835 Patent consistently uses "direct address translation" to refer to translating an address using page tables after a TLB miss. Ex. 29 at 22:6-10 ("In the event of a TLB miss, a TLB miss signal 372 is coupled to the direct address translation unit 280. The DAT 280 provides page table access as illustrated at 374, and provides replacement of TLB lines as illustrated at 375."); Fig. 9. Thus, ***the only enabled method*** of "direct address translation" relied on by the patentee during prosecution requires the use of page tables in main memory. *See id.*

Defendants' DATU construction is further supported by the claim language. Claim 8 recites a step of determining "whether translation data for the first virtual address is stored in the first translation buffer," followed by a step of "activating the direct address translation unit to translate the first virtual address when the translation data for the first virtual address is not stored in the first translation buffer." '750 Patent at claim 8. Thus, because the "direct address translation unit" is activated when there is no pre-translated address, a direct address translation must use the original source of the translation data—page tables in main memory—rather than addresses already stored in the translation buffer. *See id.* at 2:67-3:1, 8:6-9.

2. The Unit That Performs Direct Translations Must Be Shared

Although claim construction is based primarily on the intrinsic record, the Federal Circuit endorses understanding the accused product as part of the claim construction process. "[K]nowledge of [the accused] products or process provides meaningful context for … claim construction." *Wilson Sporting Goods Co. v. Hillerich & Bradsby Co.*, 442 F.3d 1322, 1326-27 (Fed. Cir. 2006) (citations omitted); *see also Scripps Clinic & Research Found. v. Genentech, Inc.*, 927 F.2d 1565, 1580 (Fed. Cir. 1991) ("the particular accused product (or process) is kept in mind, for it is efficient to focus on the construction of only the disputed elements.").

The dispute relating to DATU centers on what is being "shared by all instruction

execution pipelines." The following diagram is an excerpt from VPT's infringement contentions and highlights what VPT seeks to define as "shared" by the instruction execution pipelines. In this diagram, the pipelines are alleged to be "each core in a Cortex A9 mpcore processor" (*i.e.*, CPU0, CPU1, CPU2, CPU3 in the diagram). Ex. 32 at 2. The alleged DATU includes "the ***group*** of Memory Management Units (MMU's)" which are each contained entirely within an individual CPU (*i.e.*, alleged pipeline). *Id.* (emphasis added).

Each core in a Cortex-A9 mpcore processor contains an instruction pipeline

The Cortex-A9 processor includes the following features :

- superscalar, variable length, out-of-order pipeline with dynamic branch prediction

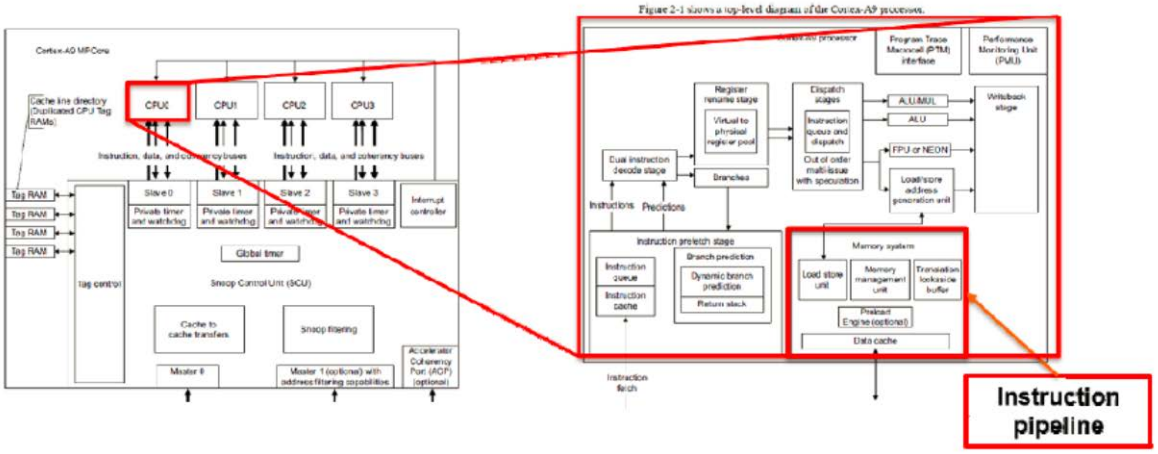

http://infocenter.arm.com/help/topic/com.arm.doc.ddi0388i/DDI0388I_cortex_a9_r4p1_trm.pdf

The *Cortex-A9 mpcore processor includes a direct address translation unit, comprised of the group of Memory Management Units ("MMUs"), the L2 cache, and the Snoop Control Unit ("SCU"). The MMUs*

Ex. 32 at 2. There is (and can be) no dispute that each "MMU" is separately contained within each separate CPU and operates independently for its associated CPU to translate data. As shown by these contentions, VPT asserts that "shared among all instruction execution pipelines" includes drawing a line around separate components contained entirely within different pipelines as one "shared" DATU. *See id.* This ambiguity is the point of VPT's vague claim construction.

VPT cannot rely on this box drawing because its arguments to PTO (1) require the unit

that directly translates addresses to be shared by multiple instruction pipelines, and (2) preclude "sharing" a collection of separate units that are distinct to each pipeline. Specifically, to distinguish prior art, VPT argued that the DATU must include shared circuitry for performing direct address translations:

> the VAX 8800 contains two processor cores. … Each processor core is substantially identical, including its own E Box, I Box, C Box and M Box. … In the event of a translation buffer miss, the respective C Box of the core for which the miss occurred executes the microcode to update its translation buffer with the correct address translation. … That is, the ***C Boxes of each core operate independently to update the respective translation buffers. The two cores do not share any "direct address translation unit,"*** …

Ex. 26 (VPT's IPR Resp.) at 6-7 (emphasis added). Thus, as VPT told the PTO, a group of separate units in "substantially identical" processing cores that "operate[] independently to update the respective translation buffer" cannot meet the DATU limitation because they do not share the same unit that performs "direct address translation"—*i.e.* translation of virtual address to physical address using data in the page tables. *See id.* Moreover, because the PTO found that Claim 8 requires the same DATU to translate both a "first virtual address" (in a first instruction pipeline) and a "second virtual address" (in a second instruction pipeline), the DATU term must refer to a common unit that performs direct address translations for both pipelines, not a collection of distinct units. *See* Ex. 33 (PTO IPR Order) at 16-18. Defendants' construction resolves the ambiguity of VPT's construction by clarifying that the DATU term requires the unit that directly translates addresses using the page tables in main memory to be shared by the pipelines.

D. "storing the translation data for the [first/second/third] virtual address from the master translation memory into the [first/second/third] translation buffer" (claims 8-12)

Plaintiff's Construction	Defendants' Construction
Plain and ordinary meaning	"storing the translation data for the [first/second/third] virtual address from the

	master translation memory into the [first/second/third] translation buffer even when the translation data for the [first/second/third] virtual address is stored in the [first/second/third] translation buffer."

The only dispute for this term is whether the "storing" step has a condition precedent. Specifically, while the '750 Patent specification discloses only updating the translation buffer when the comparator issues a miss signal—*i.e.*, only when translation data for the virtual address is not already stored in the translation buffer—Defendants' construction merely seeks to clarify that the claim language does not have any such required condition precedent.[9] *See, e.g.*, '750 Patent at 5:6-5:26; 3:40-48; 3:56-58.

Claim 8 recites a step of "indicating whether translation data for the first virtual address is stored in the first translation buffer," followed by a step of "activating the direct address translation unit to translate the first virtual address when the translation data for the first virtual address is not stored in the [first/second] translation buffer." *Id.* at claim 8. The "indicating" step indicates whether the translation data for the first virtual address was already stored in the first translation buffer (*i.e.*, a TLB hit), or whether that data is not there (*i.e.*, a TLB miss). The following "activating" step recites a condition precedent that the "activating" step only occurs if there is a TLB miss as determined by the "indicating" step. *See id.* The "storing" step comes next and is structurally and grammatically set forth as another step in the method independent of whether the "indicating" step indicated a hit or miss, and independent of whether the conditional "activating" step occurred. The "storing" step, therefore, neither recites nor requires a condition precedent for storing translation data from the master translation memory into the [first/second] translation buffer. Thus, Defendants' construction clarifies the plain meaning of the "storing …"

limitation by requiring a storing step to occur regardless of whether there is a hit or miss signal from the "indicating" step.

E. "whenever translation data for the [first/second] virtual address from the master translation memory is stored into the [first/second] translation buffer" (claims 9-10)

Plaintiff's Construction	Defendants' Construction
"whenever" should be construed to mean "at the same time when."	"whenever" should be construed to mean "every time that."

Defendants seek a construction using the plain and ordinary definition of "whenever"—which is consistent with the intrinsic record and taken verbatim from the specification. VPT did not address this term in its brief, and its proposed construction contradicts the intrinsic record.

The difference between the competing constructions is subtle, but important. Under Defendants' construction, there is a causal link between two actions: Action B occurs "whenever" Action A occurs means the performance of Action A necessarily causes Action B to occur. VPT's construction removes that linkage by allowing the limitation to be met by the (serendipitous) situation in which Action A and Action B happen to occur at the same time through independent actions. This construction misconstrues "whenever" and should be rejected. *See Mars, Inc. v. Coin Acceptors, Inc.*, 478 F. Supp. 2d 689, 711-12 (D.N.J. 2007) (construing "whenever" to require a first event to occur "every time" a second event occurs).

The claim language supports Defendants' construction. Claim 8 recites the steps of storing translation data for a first virtual address into a first translation buffer, and storing translation data for a second virtual address into a second translation buffer. *See* '750 Patent at claim 8. Dependent claim 9 adds the requirement that the first translation data is also stored into the second translation buffer "whenever translation data for the first virtual address … is stored

[9] Defendants do not concede that the "storing" terms as construed satisfy the requirements of 35

into the first translation buffer." *Id.* at claim 9. The "whenever" clause of claim 9 clearly unites and links the two actions of storing the translation data into the "first translation buffer" and the storing of the same translation data into the "second translation buffer." *See id.* Defendants' construction reflects the recited linkage between the two "storing" actions.

Defendants' construction is also supported by the specification, which discloses an embodiment that links multiple storage actions and equates "whenever" with "every time":

> update control circuit 240 can be hardware or software programmed to simultaneously update all TLB's with the same translation data ***whenever*** the address translation tables in main memory are accessed to obtain address translation information for any other TLB. That is, ***every time*** DTU 162 is activated for translating a virtual address supplied by pipeline 210A, then update control circuit stores the translation data in each of TLB's 222A-C.

Id. at 5:39-46 (emphasis added). In this embodiment, ***every time*** one instruction pipeline ("pipeline 210A") calculates translation data, the resulting translation data is stored in both its own translation buffer and the translation buffers of other instruction pipelines ("each of TLB's 222A-C"). *See id.* This avoids requiring another pipeline to independently calculate the translation data. *See id.* at 3:56-61; 4:35-37.

VPT's construction negates any linkage between the two actions and misconstrues "whenever."[10] Defendants' construction reflects the plain and ordinary meaning of "whenever" and is fully consistent with the intrinsic record.

IV. CONCLUSION

For the foregoing reasons, Defendants respectfully requests that the Court enter an order adopting Defendants' proposed claim constructions.

U.S.C. § 112, including the written description and definiteness requirements.

[10] This is confirmed by VPT's infringement contentions which merely require the two translation buffers be updated with the same translations—even if those translations are independently calculated by each instruction pipeline in the serendipitous situation that the pipelines are

Dated: December 23, 2014

Respectfully submitted,

By: */s/ Ryan K. Yagura (with permission)*
Ryan K. Yagura - Lead Attorney
Texas State Bar No. 24075933
ryagura@omm.com
Nicholas J. Whilt
Admitted Pro Hac Vice
nwhilt@omm.com
Xin-Yi Zhou
Admitted Pro Hac Vice
vzhou@omm.com
O'MELVENY & MYERS LLP
400 South Hope Street
Los Angeles, California 90071-2899
Telephone: (213) 430-6000
Facsimile: (213) 430-6407

Geoffrey Hurndall Yost
Admitted Pro Hac Vice
gyost@omm.com
O'MELVENY & MYERS LLP
Two Embarcadero Center, 28th Floor
San Francisco, California 94111-3305
Telephone: (415) 984-8700
Facsimile: (415) 984-8701

Melissa Richards Smith
GILLAM & SMITH, LLP
303 South Washington Avenue
Marshall, TX 75670
Telephone: (903) 934-8450
Facsimile: (903) 934-9257

ATTORNEYS FOR DEFENDANTS SAMSUNG ELECTRONICS CO., LTD., SAMSUNG ELECTRONICS AMERICA, INC., AND SAMSUNG TELECOMMUNICATIONS AMERICA, LLC

---

accessing the same memory areas. *See* Ex. 32 at 19-20("This [whenever clause] will occur, for example, when cores [instruction pipelines] are using the same area of memory").

/s/ Kevin P. Anderson (with permission)
Michael E. Jones
State Bar No. 10929400
mikejones@potterminton.com
POTTER MINTON
A Professional Corporation
110 N. College, Suite 500
Tyler, TX 75702
Tel: (903) 597-8311
Fax: (903) 593-0846

Kevin P. Anderson
kanderson@wileyrein.com
WILEY REIN LLP
1776 K Street, NW
Washington, D.C. 20006
Tel: 202.719.7000
Fax: 202.719.7049

ATTORNEYS FOR DEFENDANTS FREESCALE SEMICONDUCTOR, INC.; HISENSE USA CORPORATION; AND LSI CORPORATION

/s/ *Joseph T. Jakubek (with permission)*
Jennifer H. Doan
Texas Bar No. 08809050
Bent D. McCabe
Texas Bar No. 24088004
HALTOM & DOAN
Crown Executive Center
Summerhill Road, Suite 100
Texarkana, TX 75503
Telephone: (903) 255-1000
Facsimile: (903) 255-0800
Email: jdoan@haltomdoan.com
Email: bmccabe@haltomdoan.com

Joseph T. Jakubek
OR Bar No. 950326
joseph.jakubek@klarquist.com
John D. Vandenberg
OR Bar No. 893755
john.vandenberg@klarquist.com
KLARQUIST SPARKMAN, LLP
121 S.W. Salmon Street, Suite 1600
Portland, Oregon 97204
Telephone: (503) 595-5300
Facsimile: (503) 595-5301

ATTORNEYS FOR DEFENDANT AMAZON.COM, INC.

/s/ Kelly P. Chen (with permission)
JAMIL N. ALIBHAI
Texas State Bar No. 00793248
jalibhai@munckwilson.com
ROBERT D.MCCUTCHEON
Texas State Bar No. 00789480
rmccutcheon@munckwilson.com
KELLY P. CHEN
Texas State Bar No. 24062664
kchen@munckwilson.com
MUNCK WILSON MANDALA, LLP
12770 Coit Road, Ste. 600
Dallas, Texas 75251
(972) 628-3600 Telephone
(972) 628-3616 Facsimile

ATTORNEYS FOR DEFENDANT HUAWEI DEVICE USA, INC. AND FUTUREWEI TECHNOLOGIES, INC.

/s/ Michael L. Raspino (with permission)
Winstol D. Carter, Jr.
Bar No. 03932950
wcarter@morganlewis.com
Michael L. Raspino
Bar No. 24072513
mraspino@morganlewis.com
MORGAN, LEWIS & BOCKIUS LLP
1000 Louisiana Street
Suite 4000
Houston, TX 77002
T. 713.890.5000
F. 713.890.5001

Collin W. Park (*pro hac vice*)
DC Bar No. 470486
cpark@morganlewis.com
Mark W. Taylor (*pro hac vice*)
DC Bar No. 493843
mark.taylor@morganlewis.com
MORGAN, LEWIS& BOCKIUS LLP
1111 Pennsylvania Avenue, NW
Washington, DC 20004
T. 202.739.3000
F. 202.739.3001

ATTORNEYS FOR DEFENDANT LG ELECTRONICS USA, INC.

By: */s/ Amanda A. Abraham (with permission)*
Carl R. Roth
cr@rothfirm.com
Brendan C. Roth
br@rothfirm.com
Amanda A. Abraham
aa@rothfirm.com
THE ROTH LAW FIRM, P.C.
115 N. Wellington, Suite 200
Marshall, Texas 75670
Telephone: (903) 935-1665
Facsimile: (903) 935-1797

Of Counsel:
John J. Patti
j-patti@ti.com
Sarah Vollbrecht
s-vollbrecht@ti.com
Texas Instruments Incorporated
P.O. Box 655474, MS 3999
Dallas, TX 75265-5474
Telephone: (214) 479-1216

ATTORNEYS FOR DEFENDANT TEXAS INSTRUMENTS, INC.

By: */s/ Joseph M. Casino (with permission)*
Melvin ("Trip") R. Wilcox III
(State Bar No. 21454800)
YARBROUGH WILCOX GUNTER PLLC
100 E. Ferguson Street - Suite 1015
Tyler, TX 75702
Tel: 903.595.1133
Fax: 903.595.0191
mrw@yw-lawfirm.com

Joseph M. Casino (*Admitted Pro Hac*)
Abraham Kasdan (*To be Admitted Pro Hac*)
Wiggin and Dana LLP
450 Lexington Avenue
New York, New York 10017
Telephone: (212) 551-2842
Facsimile: (212) 490-0536
jcasino@wiggin.com
akasdan@wiggin.com

ATTORNEYS FOR DEFENDANT PANASONIC CORPORATION OF NORTH AMERICA

/s/ Christopher P. Broderick (with permission)
J. Thad Heartfield
Texas Bar No. 09346800
thad@heartfieldlawfirm.com
THE HEARTFIELD LAW FIRM
2195 Dowlen Road
Beaumont, Texas 77706
Telephone: 409.866.3318
Facsimile: 409.866.5789

Christopher P. Broderick
cbroderick@orrick.com
ORRICK, HERRINGTON & SUTCLIFFE
777 South Figueroa Street, Suite 3200
Los Angeles, California 90017
Telephone: 213.629.2020
Facsimile: 213.612.2499

ATTORNEYS FOR DEFENDANT SHARP ELECTRONICS CORPORATION

/s/ Christian G. Andreu-Von Euw (with permission)
Melissa Smith
GILLAM & SMITH LLP
303 South Washington Avenue
Marshall, Texas 75670
Telephone: 903.934.8450
Facsimile: 903.934.9257

M. Andrew Woodmansee
(Ca Bar No. 201780)
Mawoodmansee@Mofo.Com
Christian G. Andreu-Von Euw
(Ca Bar No. 265360)
Christian@Mofo.Com
Morrison & Foerster Llp
12531 High Bluff Drive, Suite 100
San Diego, California 92130-2040
Telephone: 858.720 5100
Facsimile: 858.720.5125

ATTORNEYS FOR DEFENDANT KYOCERA COMMUNICATIONS, INC.

/s/ William J. McCabe (with permission)
John F. Bufe
Texas State Bar No. 03316930
POTTER MINTON, PC
110 N. College Avenue
Suite 500
Tyler, TX 75702
(903)-597-8311 (telephone)
(903)-593-0846 (facsimile)
johnbufe@potterminton.com

William J. McCabe (*pro hac vice*)
Meera Nair *(pro hac vice*)
ROPES & GRAY LLP
1211 Avenue of the Americas
New York, New York 10036
(212) 596-9000 (telephone)
(212) 596-9090 (facsimile)
william.mccabe@ropesgray.com
meera.nair@ropesgray.com

Maxwell A. Fox (*pro hac vice*)
Han Xu (*pro hac vice*)
ROPES & GRAY LLP
Yusen Building 2F, 3-2 Marunouchi 2-Chome
Chiyoda-ku, Tokyo 100-0005, Japan
+81-3-6259-3500 (telephone)
+81-3-6259-3501 (facsimile)
maxwell.fox@ropesgray.com
han.xu@ropesgray.com

Jenna M. Gillingham
ROPES & GRAY LLP
191 North Wacker Drive
Floor 32
Chicago, IL 60606
(312)-845-1200 (telephone)
(312)-845-5592
jenna.gillingham@ropesgray.com

ATTORNEYS FOR DEFENDANT
NEC CORPORATION OF AMERICA

/s/ John M. Hintz (with permission)
Michael E. Jones
State Bar No. 10929400
Potter Minton
110 N. College, Suite 500
Tyler, Texas 75702
Tel.: (903) 597-8311
Fax: (903) 593-0846

John M. Hintz (admitted in E.D. Texas)
john.hintz@rimonlaw.com
RIMON, P.C.
245 Park Avenue – 39th Floor
New York, NY 10167
Tel: 917.633.4913

ATTORNEYS FOR MEDIATEK USA INC.

CERTIFICATE OF SERVICE

The undersigned certifies that the foregoing document was filed electronically in compliance with Local Rule CV-5(a). As such, this document was served on all counsel who are deemed to have consented to electronic service. Local Rule CV-5(a)(3)(A). Pursuant to Fed. R. Civ. P. 5(d) and Local Rule CV-5(d) and (e), all other counsel of record not deemed to have consented to electronic service were served with a true and correct copy of the foregoing by email, on this 23rd day of December, 2014.

/s/ Melissa Smith
Melissa Smith